arXiv:1411.5172v1 [cs.LG] 19 Nov 2014

Learning nonparametric differential equations with operator-valued kernels and gradient matching

Markus Heinonen[1,2], Florence d'Alché-Buc[3,2]

[1] IRSN, PRP-HOM/SRBE/L3R, 31 avenue de la Division Leclerc 92262 Fontenay-aux-roses, France

[2] IBISC, Universit d'Evry, 23 BD de France, Evry 91037 cedex France

[3] LTCI UMR cnrs 5141, Institut Mines-Télécom, Télécom ParisTech, 37 rue Dareau - 75014 Paris, France

November 20, 2014

Abstract

Modeling dynamical systems with ordinary differential equations implies a mechanistic view of the process underlying the dynamics. However in many cases, this knowledge is not available. To overcome this issue, we introduce a general framework for nonparametric ODE models using penalized regression in Reproducing Kernel Hilbert Spaces (RKHS) based on operator-valued kernels. Moreover, we extend the scope of gradient matching approaches to nonparametric ODE. A smooth estimate of the solution ODE is built to provide an approximation of the derivative of the ODE solution which is in turn used to learn the nonparametric ODE model. This approach benefits from the flexibility of penalized regression in RKHS allowing for ridge or (structured) sparse regression as well. Very good results are shown on 3 different ODE systems.

1 Introduction

Dynamical systems modeling is a cornerstone of experimental sciences. In biology, as well as in physics and chemistry, modelers attempt to capture the dynamical behavior of a given system or a phenomenon in order to improve its understanding and eventually make predictions about its future state. Systems of coupled ordinary differential equations (ODEs) are undoubtedly the most widely used models in science. A single ordinary differential equation describes the change in rate of a given state variable as a function of other state variables. A set of coupled ODEs together with some initial condition can account for the full dynamics of a system. Even simple ODE functions can describe complex dynamical behaviours (Hirsch et al., 2004).

Ode modeling consists of the two tasks of (i) choosing the parametric ODE model, which is usually based on the knowledge of the system under study, and (ii) estimation of the parameters of the system from noisy observations. In many cases there is no obvious choice in favor of a given model. If many models are in competition, the choice can rely on hypothesis testing based model selection (Cox, 1961; Vuong, 1989). However it is not always easy to fulfill the assumptions of the statistical tests or propose a single mechanistic model. The aim of this work is to overcome these issuees by introducing a radically new angle for ODE modeling using nonparametric models, which sidestep the issue of model choice and provide a principled approach to parameter learning. A nonparametric model does not necessitate prior knowledge of the system at hand.

More precisely, we consider the dynamics of a system governed by the following first-order multivariate ordinary differential equations:

$$\dot{\mathbf{x}}(t) = f_{true}(\mathbf{x}(t)), \tag{1}$$

$$\mathbf{x}(t_0) = x_0, \tag{2}$$

where $\mathbf{x}(t) \in \mathbb{R}^p$ is the state vector of a p-dimensional dynamical system at time t, e.g. the ODE solution given the initial condition $\mathbf{x}_0$, and the $\dot{\mathbf{x}}(t) = \frac{d\mathbf{x}(t)}{dt}$ is a first order derivative of $\mathbf{x}(t)$ over time and f_{true} is a vector-valued

function. The ODE solution satisfies

$$\mathbf{x}(t) = \mathbf{x}_0 + \int_0^t f_{true}(\mathbf{x}(\tau))d\tau.$$

We assume that f_{true} is unknown and we observe an n-length multivariate time series $(\mathbf{y}_0, \ldots, \mathbf{y}_{n-1})$ obtained from an additive noisy observation model at discrete time points $t_0, \ldots, t_{n-1}$:

$$\mathbf{y}_\ell = \mathbf{x}(t_\ell) + \epsilon_\ell, \tag{3}$$

where ϵ_ℓ's are i.i.d samples from a Gaussian distribution.

In classical methods ODE approaches the parameters θ of the function $f(\mathbf{x}(t), \theta)$ are estimated with least squares approach, where the solution is simulated using a trial set of initial values for θ, with subsequent parameter optimisation to maximise the simulated solution $\hat{x}_\theta(t_\ell) = \int_0^{t_\ell} f(\mathbf{x}(\tau), \theta)d\tau$ match against the observations $\mathbf{y}_\ell$. This approach, classically proposed off-the-shelf, involves computationally-intensive approximations and suffers from the scarcity of the observations. To overcome these issues, a family of methods proposed under different names of *collocation methods* (Varah, 1982), *gradient matching* (Ellner et al., 2002) or *profiled estimation* (Ramsay et al., 2007) have proposed to produce a smooth approximation $\hat{\mathbf{x}}(t)$ that can be used in turn as a surrogate $\dot{\hat{\mathbf{x}}}(t)$ of the true derivative $\dot{\mathbf{x}}(t)$ (Bellman and Roth, 1971). Then, we optimise the match between $\dot{\hat{\mathbf{x}}}$ and $f(\hat{\mathbf{x}}, \theta)$, which does not require the costly integration step. In iterated estimation procedure the smoother and parameter estimation are iterated to correct the approximation errors in both terms. Using just two-steps has been analyzed in the parametric case in terms of asymptotics and it enjoys consistency results and provides nearly equal performance in the finite sample case (Brunel, 2008; Gugushvili and Klaassen, 2011).

In this work, we adopt the two-step gradient matching approach to the nonparametric ODE estimation to learn a nonparametric model $h \in \mathcal{H}$ to estimate the function f_{true}:

Step 1. Learn $g : \mathbb{R} \to \mathbb{R}^p$ in a functional space $\mathcal{G}$ to approximate the observed trajectory $\mathbf{y}_0, \ldots, \mathbf{y}_{n-1}$ at time points $t_0, \ldots, t_{n-1}$ with $t_0 = 0$.

Step 2. Given an approximation g of the observed trajectory, learn $h : \mathbb{R}^p \to \mathbb{R}^p$ in a functional space $\mathcal{H}$ to approximate the differential $\dot{g}(t)$ with $h(g(t))$.

The two steps play a very different role: with g, we want to capture the observed trajectories with enough accuracy so that $g(t)$ is close to the true value $x(t)$, and so that the derivative $\dot{g}(t)$ is close to the true differential $\dot{x}(t)$. When the covariance noise is assumed to be isotropic, each coordinate function $g_j; j = 1, \ldots, p$ of vector-valued function g can be learned separately, which we assume in this paper. In the second step, we want to discover dependencies between the vector spaces of state variables and state variable derivatives, and for that reason we will turn to vector-valued function approximation. For that purpose, we propose as a key contribution to use the general framework of penalized regression for vector-valued functions in Reproducing Kernel Hilbert Spaces (RKHS) extending our preliminary works on the subject (d'Alché Buc et al., 2012; Lim et al., 2013, 2014). The RKHS theory offers a flexible framework for regularization in functional spaces defined by a positive definite kernel, with well-studied applications such in scalar-valued function approximation, such as the SVM (Aronszajn, 1950; Wahba, 1990). Generalising the RKHS theory to vector-valued functions with operator-valued kernels (Pedrick, 1957; Senkene and Tempel'man, 1973; Micchelli and Pontil, 2005; Caponnetto et al., 2008) has recently attracted a surge of attention for vector-valued functions approximation. Operator-valued kernels provide an elegant and powerful approach for multi-task regression (Micchelli and Pontil, 2005), structured supervised and semi-supervised output prediction (Brouard et al., 2011; Dinuzzo, 2011), functional regression (Kadri et al., 2011) and nonparametric vector autoregression (Lim et al., 2013).

We introduce non-parametric operator-valued kernel-based regression models for gradient matching under ℓ_2 penalties in Section 2. Section 3 extends the framework to learning from multiple time series obtained from multiple initial conditions. Section 4 proposes methods to introduce sparsity to the ODE models with ℓ_1 and ℓ_1/ℓ_2 norms. Section 5 presents a method to learn the kernel function. Section 6 highlights the performance of the approach in two case studies of non-trivial classic ODE models and a realistic biological dataset with unknown model structure. We conclude in Section 7.

2 Gradient Matching for nonparametric ODE

We consider the following loss function based on empirical losses and penalty terms:

$$\mathcal{L}(h,g) = \frac{1}{2}\sum_{\ell=0}^{n-1} ||g(t_\ell) - \mathbf{y}_\ell||^2 + \lambda_g ||g||^2 + \frac{1}{2}\sum_{\ell=1}^{m} ||\dot{g}(t_\ell) - h(g(t_\ell))||^2 + \lambda_h ||h||^2. \tag{4}$$

We present how **Step 1** is solved using the classic tools of penalized regression in RKHS of scalar-valued functions and then introduce the tools based on operator-valued kernels devoted to vector-valued functions.

2.1 Learning the smoother g

Various methods for data smoothing have been proposed in the literature and most of them (such as splines) can be described in the context of functional approximation in Reproducing Kernel Hilbert Spaces(see the seminal work of Wahba (1990) and also Pearce and Wand (2006)). We apply standard kernel ridge regression as such an approximation.

For each state variable indexed by $j = 1, \ldots, p$, we choose a positive definite kernel $k_j : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ and define the Hilbert space $\mathcal{G}_j$. We build a smoother

$$g_j(t) = \sum_{i=0}^{n-1} b_{ij} k_j(t, t_i)$$

in the Hilbert space $\mathcal{G}_j$ by solving a kernel ridge regression problem. Given the observed data $\mathcal{S}_j = \{(t_0, y_{0j}), \ldots, (t_{n-1}, y_{(n-1)j})\}$, minimizing the following loss:

$$\mathcal{L}(g_j; \mathcal{S}_j) = \frac{1}{2}\sum_{i=0}^{n-1} (y_{ij} - g_j(t_i))^2 + \lambda_{g,j} \left\|g_j\right\|^2, \tag{5}$$

with $||g_j||^2 = \mathbf{b}_j^T K_j \mathbf{b}_j$ leads to a unique minimizer:

$$\mathbf{b}_j = (K_j + \lambda_{g,j} Id)^{-1} \mathbf{y}_j, \tag{6}$$

where $\mathbf{b}_j$ is the n-dimensional parameter vector of the solution model g_j (for sake of simplicity, we avoid the hat notation), K_j is the Gram kernel matrix computed on input data $t_0, \ldots, t_{n-1}$, Id is the $n \times n$ identity matrix, $\mathbf{y}_j$ is n-dimensional column vector of observed variable j.

The derivatives $\dot{g}(t) = (\dot{g}_1(t), \ldots, \dot{g}_p(t))$ are straightforward to calculate as

$$\dot{g}_j(t) = \sum_{i=0}^{n-1} b_{ji} \frac{dk_j(t, t_i)}{dt}. \tag{7}$$

2.2 Learning h with operator-valued kernels

Operator-valued kernel (Álvarez et al., 2012) extends the well-known scalar-valued kernels in order to deal with functions with values in Hilbert Spaces. We present briefly the fundamentals of operator-valued kernels and the associated RKHS theory as introduced in (Senkene and Tempel'man, 1973; Micchelli and Pontil, 2005). Then we apply this theory to the case of functions with values in $\mathbb{R}^p$.

Let $\mathcal{X}$ be a non-empty set and $\mathcal{Y}$ a Hilbert space. We note $\mathcal{L}(\mathcal{Y})$, the set of all bounded linear operators from $\mathcal{Y}$ to itself. Given $M \in \mathcal{L}(\mathcal{Y})$, M^* denotes the adjoint of M.

Definition 1 (Operator-valued kernel)**.** *Let $\mathcal{X}$ be a non-empty set. $K : \mathcal{X} \times \mathcal{X} \to \mathcal{L}(\mathcal{Y})$ is an operator-valued kernel if:*

- $\forall (x, x') \in \mathcal{X} \times \mathcal{X}$, $K(x, x') = K(x', x)^*$

- $\forall m \in \mathbb{N}$, $\forall \mathcal{S}_m = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^m \subseteq \mathcal{X} \times \mathcal{Y}$,

$$\sum_{i,j=1}^{m} \langle \mathbf{y}_i, K(\mathbf{x}_i, \mathbf{x}_j)\mathbf{y}_j \rangle_{\mathbb{R}^p} \geq 0.$$

Similarly to the scalar case, an operator-valued kernel allows to build a unique Reproducing Kernel Hilbert Space $\mathcal{H}_K$. First, the span of $\{h(\cdot) = \sum_i K(\cdot, x_i)\mathbf{a}_i, x_i \in \mathcal{X}\}$) is endowed with the following inner product:

$$\langle f, g \rangle_{\mathcal{H}_K} = \sum_{i,j} \langle \mathbf{a}_i, K(x_i, z_j)\mathbf{b}_j \rangle$$

with $f = \sum_i K(\cdot, x_i)\mathbf{a}_i$ and $g = \sum_j K(\cdot, z_j)\mathbf{b}_j$. This choices ensures the reproducing property:

$$\forall (x, \mathbf{y}, f) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{H}, \langle f, K(\cdot, x)\mathbf{y} \rangle_{\mathcal{H}_K} = \langle f(x), \mathbf{y} \rangle_{\mathcal{Y}}$$

Then the corresponding norm $\| \cdot \|_{\mathcal{H}_0}$ is defined by $\| f \|^2_{\mathcal{H}_K} = \langle f, f \rangle_{\mathcal{H}_0}$. Then $\mathcal{H}$ is completed by including the limits of Cauchy sequences for which the reproducing property still holds. Approximation of a function in a RKHS enjoys representer theorems such as the following general one proved by Micchelli and Pontil (2005)

Theorem 1 (Micchelli and Pontil (2005))**.** *Let $\mathcal{X}$ be a non-empty set, $\mathcal{Y}$ a Hilbert Space and Kan operator-valued kernel with values in $\mathcal{L}(\mathcal{Y})$. Let $\mathcal{H}$ the RKHS built from K. Let $\{(\mathbf{x}_\ell, \mathbf{y}_\ell) \in \mathcal{X} \times \mathcal{Y}, \ell = 1, \ldots, m\}$, a given set. Let $V : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ be a loss function, and $\lambda_h > 0$ a regularization parameter. Then any function inn $\mathcal{H}$ minimizing the following cost function:*

$$\mathcal{J}(h) = \sum_{\ell=1}^{m} V(h(x_\ell), \mathbf{y}_\ell) + \lambda_h \|h\|^2_{\mathcal{H}} \ ,$$

admits an expansion:

$$h(\cdot) = \sum_{\ell=1}^{m} K(\cdot, x_\ell)\mathbf{a}_\ell \ ,$$

where the coefficients $\mathbf{a}_\ell, \ell = 1, \ldots, m$ are vectors in the Hilbert space $\mathcal{Y}$.

To solve the **Step 2** given g, we want to find a function

$$h(\mathbf{x}) = \sum_{\ell=1}^{m} K(\mathbf{x}, g(t_\ell))\mathbf{a}_\ell \tag{8}$$

intuitively minimizing the expected square gradient matching error over the time interval of interest $(0, T)$, plus a regularising term:

$$\mathcal{L}(h \,|\, g) = \frac{1}{2} \int_0^T ||\dot{g}(t) - h(g(t))||^2 dt + \lambda_h ||h||^2.$$

However, to apply the representer theorem to h, we are inclined to replace this expectation by an empirical mean (ignoring a factor) as

$$\mathcal{L}(h \,|\, g, \tau_1, \ldots, \tau_m) = \frac{1}{2} \sum_{\ell=1}^{m} ||\dot{g}(\tau_\ell) - h(g(\tau_\ell))||^2 + \lambda_h ||h||^2 \tag{9}$$

with $\tau_1, \ldots, \tau_m$ a sequence of m positive reals, uniformly and independently sampled. We effectively use m values along the trajectory estimate g to act as the dataset to learn h from, where usually $m \geq n$. Now, the representer theorem 1 applies to h with the following choices: $\mathcal{X} = \mathcal{Y} = \mathbb{R}^p$, $\mathbf{x}_\ell = g(\tau_\ell), \ell = 1, \ldots, m$ and $\mathbf{y}_\ell = \dot{g}(\tau_\ell), \ell = 1, \ldots, m$. This can be re-formulated as the Gradient Matching Representer Theorem:

Theorem 2 (Gradient Representer Theorem)**.** *Let $g : \mathbb{R} \to \mathbb{R}^p$ be a vector-valued function, differentiable and $\tau_\ell, \ell = 1, \ldots, m$, a set of positive reals. Let $K : \mathbb{R}^p \times \mathbb{R}^p \to \mathbb{R}$ a positive matrix-valued kernel and $\mathcal{H}_K = \mathcal{H}$ the RKHS built from K. Let $V : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ be a loss function, and $\lambda_1 > 0$ a regularization parameter. Then any function of $\mathcal{H}$ minimizing the following cost function:*

$$\mathcal{J}(h) = \sum_{\ell=1}^{m} V(h(g(\tau_\ell)), \dot{g}(\tau_\ell)) + \lambda_1 \|h\|_{\mathcal{H}}^2,$$

admits an expansion

$$h(\cdot) = \sum_{\ell=1}^{m} K(\cdot, g(\tau_\ell))\mathbf{a}_\ell,$$

where the coefficients $\mathbf{a}_\ell, \ell = 1, \ldots, m$ are vectors in $\mathbb{R}^p$.

Step 2 is then solved as $\min_{h \in \mathcal{H}} \mathcal{L}(h \mid g, \tau_1, \ldots, \tau_m)$, given the kernel ridge regression solution $g(t)^T = (g_1(t), \ldots, g_p(t))$. The gradient matching problem admits a closed-form solution

$$\mathbf{a} = (\mathbf{K} + \lambda_h Id)^{-1}\dot{\mathbf{g}},$$

where the pm-dimensional $\dot{\mathbf{g}}$ is obtained by stacking the column vectors $\dot{g}(\tau_1), \ldots, \dot{g}(\tau_m)$, and $\mathbf{a}$ is a pm-dimensional vector obtained by stacking the column vectors $\mathbf{a}_1, \ldots, \mathbf{a}_m$. $\mathbf{K}$ is a $m \times m$ block matrix, where each block is a matrix of size $p \times p$. The (ℓ, s)'th block of $\mathbf{K}$ corresponds to the matrix $K(g(\tau_\ell), g(\tau_s))$.

2.3 Operator-valued kernel families

Several operator-valued kernels have been defined in the literature (Micchelli and Pontil (2005); Álvarez et al. (2012); Lim et al. (2013)). We use two such kernels, which both are universal, SDP and based on the gaussian scalar kernel. First, the *decomposable* kernel

$$K_{DC}(\mathbf{x}, \mathbf{x}') = k(\mathbf{x}, \mathbf{x}')C,$$

where $k(\cdot, \cdot)$ is a standard gaussian scalar kernel, and $C \in \mathbb{R}^{p \times p}$ is a positive definite dependency structure matrix. Second, the *transformable* kernel

$$K_{TF}(\mathbf{x}, \mathbf{x}')_{ij} = k(\mathbf{x}^i, \mathbf{x}'^j)$$

measures the pairwise similarities between features of data points, i.e. the variables between state vectors. Finally, we can also use the *Hadamard* kernel

$$K_{had}(\mathbf{x}, \mathbf{x}') = K_{DC} \circ K_{TF} = k(\mathbf{x}, \mathbf{x}')C \circ K_{TF}(\mathbf{x}, \mathbf{x}').$$

3 Learning from multiple initial conditions

In parametric ODEs, it is well known that using time series coming from different initial conditions reduces the non-identifiability of parameters. Similarly we also want to increase the accuracy of our estimate h with multiple time series in the nonparametric case. However, in contrast to the parametric case, using r multiple datasets produces r models h^i ($i = 1, \ldots, r$). Given the assumption that there exists a true ODE model, it is supposed to be unique. We therefore propose to learn r nonparametric models h^i with a smoothness constraint that imposes that they should be close in terms of ℓ_2 norm in the functional space $\mathcal{H}$, and hence they should give similar estimates for the same input. This corresponds to a multi-task approach (Evgeniou and Pontil, 2004) and is strongly related to the recent general framework of manifold regularization (Ha Quang and Bazzani, 2013).

Let us assume that r multivariate time series are observed, starting from r different initial conditions. For simplicity, we assume that each time serie has the same length. **Step 1** consists of learning r vector-valued functions

$g^i, i = 1, \ldots, r$ from $\mathcal{S}_n^i = \{(t_0, \mathbf{y}_0^i), \ldots, (t_{n-1}, \mathbf{y}_{n-1}^i)\}$ as described in Section 2. The new loss to be minimized is

$$\begin{aligned}\mathcal{L}(h^1, \ldots, h^r | g^1, \ldots, g^r) &= \sum_{i=1}^{r} \mathcal{L}(h^i | g^i) + \frac{1}{4}\lambda \sum_{i,j=1}^{r} ||h^i - h^j||^2 \qquad (10)\\ &= \frac{1}{2}\sum_{i=1}^{r}\left(||\dot{\mathbf{g}}^i - \mathbf{K}^{ii}\mathbf{a}^i||^2 + \lambda_h {\mathbf{a}^i}^T \mathbf{K}^{ii}\mathbf{a}^i\right) + \frac{1}{2}\lambda_{sim}\left(r\sum_{i=1}^{r}{\mathbf{a}^i}^T\mathbf{K}^{ii}\mathbf{a}^i - \sum_{i,j=1}^{r}{\mathbf{a}^i}^T\mathbf{K}^{ij}\mathbf{a}^j\right)\\ &= \frac{1}{2}\left\|\dot{\mathbf{g}} - \mathrm{diag}\mathcal{K}\mathbf{a}\right\|^2 + \frac{1}{2}\lambda_h \mathbf{a}^T \mathrm{diag}\mathcal{K}\mathbf{a} + \frac{1}{2}\lambda_{sim}\mathbf{a}^T(r\mathrm{diag}\mathcal{K} - \mathcal{K})\mathbf{a},\end{aligned}$$

where $\dot{\mathbf{g}}^i$ is the stacked column vectors of all $\dot{g}(\tau_\ell)^i$'s, operator-valued kernel matrices $\mathbf{K}^{ij} = (K(g^i(\tau_j), g^j(\tau_k)))_{j,k=1}^m$ are the kernel matrices comparing time-series i and j, and $\mathbf{a}^i$ is the stacked column vectors $(\mathbf{a}_1^i, \ldots, \mathbf{a}_m^i)$. Furthermore, the $\dot{\mathbf{g}}$ is concatenation of all $\dot{\mathbf{g}}^i$'s into a single column matrix, $\mathbf{a}$ is a concatenation of all $\mathbf{a}^i$'s, $\mathcal{K} = (\mathbf{K}^{ij})_{i,j=1}^r$ is a $r \times r$ block matrix of operator-valued kernel matrices $\mathbf{K}^{ij}$, and the diagonal $\mathrm{diag}\mathcal{K}$ has the diagonal blocks $\mathbf{K}^{ii}$ and zeroes elsewhere.

Similarly to standard kernel ridge regression and to semi-supervised kernel ridge regression, vector $\mathbf{a}$ can be obtained by annealing the gradient $\frac{\partial \mathcal{L}(h^1,\ldots,h^r|g^1,\ldots,g^r)}{\partial \mathbf{a}}$ which gives the following closed form:

$$\mathbf{a} = \left(\mathrm{diag}\mathcal{K} + (\lambda_h + r\lambda)Id - \lambda_{sim}(\mathrm{diag}\mathcal{K})^{-1}\mathcal{K}\right)^{-1}\dot{\mathbf{g}},$$

with Id being here the identity matrix of dimension mr. However to get an efficient approximation avoiding numerical issues, we use a stochastic averaged gradient descent in numerical experiments.

A single function can be constructed as the empirical average $\bar{h}(\cdot) = \frac{1}{r}\sum_{i=1}^{r} h^i(\cdot)$, representing the consensus model.

4 Extension to sparse models

When learning the function h, we can use in principle as many training samples as we wish since the function g and its analytical derivative $\dot{g}$ are available. The ℓ_2 penalty associated with $h(\cdot)$ ensures the smoothness of the estimated functions. If we use the term *support vectors* to refer to training vectors $g(\tau_\ell)$ that have a non-zero contribution to the model, then an interesting goal is to try to use as few as possible training vectors and thus to try to reduce the number of corresponding non-zero parameters of $\mathbf{a}$. To achieve this, similarly to matrix-valued kernel-based autoregressive models (Lim et al., 2014), we add to $\mathcal{L}(h|\hat{g})$ the following penalty

$$\Omega_{\mathbf{a}}(\mathbf{a}) = \Omega_1(\mathbf{a}) + \Omega_{1,2}(\mathbf{a})$$

composed of two sparsity-inducing terms:

$$\Omega_1(\mathbf{a}) = \alpha \left\|\mathbf{a}\right\|_1 \qquad (11)$$

$$\Omega_{1,2}(\mathbf{a}) = (1-\alpha)\sum_{\ell=0}^{N}\left\|\mathbf{a}_\ell\right\|_2 \qquad (12)$$

The first term $\Omega_1(\mathbf{a})$ imposes the general lasso sparsity of the estimated function by aiming to set coefficients of the concatenated vector $\mathbf{a}$ to zero without taking into account the vector structure. The second term $\Omega_{1,2}$ imposes sparsity on the number of support vectors $\mathbf{a}_\ell$. The $\Omega_{1,2}$ is called mixed ℓ_1/ℓ_2-norm or the *group lasso* (Yuan and Lin, 2006), which exhibits some interesting features: it behaves like an ℓ_1-norm over the vectors $\mathbf{a}_\ell$ while within each vector $\mathbf{a}_\ell$, the coefficients are subject to an ℓ_2-norm constraint. The term $\alpha \in [0, 1]$ gives a convex combination of ℓ_1 and group lasso penalties ($\alpha = 0$ gives group lasso, $\alpha = 1$ gives ℓ_1 penalty), while the λ_h defines the overall regularisation effect. The new loss function $\mathcal{L}(\mathbf{a}) = \mathcal{L}(h|g) + \lambda_1\Omega_{1,2}(\mathbf{a})$ is still convex and can be decomposed into two terms:

Algorithm 1 Proximal gradient descent for minimizing (13)

Inputs : Initial solution $\mathbf{a}^{(0)} \in \mathbb{R}^{np}$, Lipschitz L
Initialize : $t^{(0)} = 1$
until convergence:
1: Compute $G = \nabla_{\mathbf{a}} \mathcal{L}_s(\mathbf{a}^{(m)})$
2: $\mathbf{a}^{(m)} = \text{prox}_{\frac{\lambda}{L_{\mathbf{a}}},\alpha}(\ell_1 + \ell_1/\ell_2)\left(\mathbf{a}^{(m)} - \frac{1}{L_{\mathbf{a}}} G\right)$
3: $t^{(m+1)} = \frac{1+\sqrt{1+4t^{(m)^2}}}{2}$
4: $\mathbf{a}^{(m)} = \mathbf{a}^{(m)} + \frac{t^{(m)}-1}{t^{(m+1)}}\left(\mathbf{a}^{(m)} - \mathbf{a}^{(m-1)}\right)$

$\mathcal{L}_s(\mathbf{a})$ which is smooth and differentiable with respect to $\mathbf{a}$ and $\mathcal{L}_{ns}(\mathbf{a})$ which is non-smooth, but nevertheless convex and subdifferentiable with respect to $\mathbf{a}$:

$$\mathcal{L}(\mathbf{a}) = \mathcal{L}_s(\mathbf{a}) + \mathcal{L}_{ns}(\mathbf{a}), \tag{13}$$

where

$$\mathcal{L}_s(\mathbf{a}) = \frac{1}{2}\sum_{\ell=0}^{n-1} \|\dot{g}(\tau_\ell) - h(g(\tau_\ell))\|_2^2 + \lambda_h ||h||_K^2$$

$$\mathcal{L}_{ns}(\mathbf{a}) = \lambda_1 \alpha ||\mathbf{a}||_1 + \lambda_1(1-\alpha)\sum_{\ell=0}^{n-1} \|\mathbf{a}_\ell\|_2 .$$

Recently, proximal gradient algorithms (Combettes and Pesquet, 2011) have been proposed for solving problems of form (13) and shown to be successful in a number of learning tasks with non-smooth constraints. The main idea relies on using the proximal operator on the gradient term. For that purpose, we introduce the following notations: $L_{\mathbf{a}}$ is a Lipschitz constant – the supremum – of the gradient $\nabla_{\mathbf{a}} \mathcal{L}_s(\mathbf{a})$. For $s > 0$, the proximal operator of a function $\mathcal{L}_{ns}$ applied to some $\mathbf{v} \in \mathbb{R}^{np}$ is given by:

$$\text{prox}_s(\mathcal{L}_{ns})(\mathbf{v}) = \arg\min_{\mathbf{u}} \left\{ \mathcal{L}_{ns}(\mathbf{u}) + \frac{1}{2s}||\mathbf{u} - \mathbf{v}||^2 \right\}$$

The proximal gradient descent algorithm is presented in 1. Intermediary variables $t^{(m)}$ and $\mathbf{a}^{(m)}$ in Step 2 and Step 3 respectively are introduced to accelerate the proximal gradient method (Beck and Teboulle, 2010).

For a given vector $\mathbf{v} \in \mathbb{R}^{np}$, the proximal operator of ℓ_1 is the element-wise shrinkage or soft-thresholding operator $\text{prox}_{\mu\ell_1} : \mathbb{R}^{np} \to \mathbb{R}^{np}$

$$\text{prox}_{\mu\ell_1}(\mathbf{u}) = \left(1 - \frac{\mu}{||\mathbf{u}||_2}\right)_+ \mathbf{u},$$

while the proximal operator of the ℓ_1/ℓ_2 is the group-wise shrinkage operator $\text{prox}_{\mu}(\ell_1/\ell_2) : \mathbb{R}^{np} \to \mathbb{R}^{np}$

$$\text{prox}_{\mu\ell_1/\ell_2}(\mathbf{u})_{\mathcal{I}} = \left(1 - \frac{\mu}{||\mathbf{u}_{\mathbf{I}}||_2}\right)_+ \mathbf{u}_{\mathbf{I}},$$

where $\mathbf{u}_{\mathcal{I}}$ denotes the coefficients of $\mathbf{u}$ indexed by $\mathbf{I}$. $\mathbf{I}$ denotes the group indexes. The proximal operator of the combined ℓ_1 and ℓ_1/ℓ_2 regularisers is the so called 'sparse group lasso', and it's defined as

$$\text{prox}_{\lambda,\alpha\ell_1+\ell_1/\ell_2}(\mathbf{u}) = \text{prox}_{\lambda\alpha}(\ell_1/\ell_2)\left(\text{prox}_{\lambda(1-\alpha)}(\ell_1)(\mathbf{u})\right)$$

We initialise the proximal algorithm with the solution on the smooth part $\mathcal{L}_s(\mathbf{a})$. The gradient of the smooth part is $\nabla_{\mathbf{a}} \mathcal{L}_s(\mathbf{a}) = (\mathbf{K} + \lambda_h Id)\mathbf{a} + \dot{\mathbf{g}}$. By definition

$$||\nabla \mathcal{L}_s(\mathbf{a}) - \nabla \mathcal{L}_s(\mathbf{a}')|| \leq ||\mathbf{K} + \lambda_h I||_F \cdot ||\mathbf{a} - \mathbf{a}'||$$

giving a Lipschitz constant $L_{\mathbf{a}} = ||\mathbf{K} + \lambda_h Id||_F$.

Algorithm 2 Incremental proximal gradient descent for minimizing (15)

Inputs: Initial solution $C_0 \in \mathbb{R}^{p\times p}$, Lipschitz L_C
Initialize: $C = C_0$
until convergence:
Step 1: $C = C - \frac{1}{L_C}\nabla_C \mathcal{L}_s(C)$
Step 2: $C = \text{prox}_{\mathcal{S}_p^+}(C)$

5 Kernel learning

In order to build the vector-valued function h, various matrix-valued kernels can be chosen either among those already described in the literature (see for instance Caponnetto et al. (2008) and Álvarez et al. (2012)) or built using the closure properties of operator-valued kernels as in Lim et al. (2013). Here, as an example, we choose the decomposable kernel which was originally proposed to tackle multi-task learning problems in Micchelli and Pontil (2005). In the case of matrix-valued kernel with values in $\mathcal{L}(\mathbb{R}^p)$, the decomposable kernel $K : \mathbb{R}^p \times \mathbb{R}^p \to \mathcal{L}(\mathbb{R}^p)$ defined as follows:

$$\forall(\mathbf{x}, \mathbf{z}), K(\mathbf{x}, \mathbf{z}) = Ck(\mathbf{x}, \mathbf{z}), \tag{14}$$

where C is a $p \times p$ positive semi-definite matrix.

The non-zero elements C_{ij} reflect dependency relationship between variables. Hence, we desire to learn the matrix C of the decomposable kernel $K(x, x') = C \cdot k(x, x')$ of the ode function $h(g(t)) = \sum_{\ell=1}^{N} K_C(g(t), \mathbf{y}_\ell)\mathbf{a}_\ell$. We resort to two-step optimization where we alternatively optimize (i) the loss function (15) given a constant C, and (ii) the loss function given a constant $\mathbf{a}$ and $\mathbf{b}$:

$$\mathcal{L}(C) = \underbrace{\frac{1}{2}\sum_{\ell=1}^{m} ||\dot{g}(t_\ell) - h_C(g(t_\ell))||^2 + \lambda_h ||h||^2}_{\mathcal{L}_s(C)} + \mathbf{1}_{\mathcal{S}_p^+}(C), \tag{15}$$

where $\mathcal{S}_p^+$ is the cone of positive semidefinite $p \times p$ matrices, and $\mathbf{1}_{\mathcal{S}_p^+}$ denotes an indicator function with $\mathbf{1}_{\mathcal{S}_p^+}(C) = 0$ if $C \in \mathcal{S}_p^+$ and ∞ otherwise.

The first two terms $\mathcal{L}_s$ of the loss function are differentiable with respect to C, while the remaining terms are non-smooth, but convex and sub-differentiable with respect to C. To optimise the loss function, we employ proximal algorithms, which here reduces to projected gradient descent where after each iteration we project the value of C to the cone of SDP matrices (Richard et al., 2012). Our algorithm is presented in Algorithm 2.

The projection onto the cone of positive semidefinite matrices is

$$\text{prox}_{\mathcal{S}_p^+}(C) = \underset{X\in\mathcal{S}_p^+}{\arg\min} ||X - C||_F = U\Lambda_+ U^T,$$

where U is the orthonormal eigenvectors of C and $\Lambda_+ = \max(\Lambda, 0)$ is the non-negative eigenvalue matrix of the corresponding eigenvalues Λ.

The gradient of the smooth part of the loss function is

$$\nabla_C \mathcal{L}_s(C) = -EKA^T + AKA^T$$

where $K = (k(g(t_\ell), g(t_s)))_{\ell,s=1}^m$ is a scalar kernel matrix, $E = (\dot{g}(t_\ell) - h(g(t_\ell))_{\ell=1}^m$ is a matrix of differences and $A = (\mathbf{a}_1, \ldots, \mathbf{a}_m) \in \mathbb{R}^{p\times m}$ is a matrix with $\mathbf{a}_\ell$ as columns. Due to the SDP constraint, the gradient matrix is symmetric.

The Lipschitz constant $L_C = ||H_I K A^T||_F^2$ is due to the property

$$||\nabla_C \mathcal{L}_s(C) - \nabla_C \mathcal{L}_s(C')||_F^2 \leq ||C - C'||_F^2 \cdot ||H_I K A^T||_F^2,$$

where $H_I = (h_I(g(t_\ell)))_{\ell=1}^m$ is the matrix of predictions from $h(g(t))$. We have the linear property $H_C = CH_I$.

6 Numerical results

We perform numerical experiments on three datasets to highlight the performance of the novel OKODE framework in ODE estimation, analyze the resulting models, and finally compare our approach to the state-of-the-art ODE estimation methods.

6.1 Experimental setting

Throughout the experiments we use independent kernel ridge regression models for the variables of $g_i(t)$ with a gaussian scalar kernel $k_i(x, x') = \exp(-\gamma_i||x - x'||^2)$ with hyperparameter λ_i. For the learning of h we use a decomposable kernel $K(g(t), g(t_\ell)) = C \cdot k(g(t), g(t_\ell))$ with a gaussian scalar kernel $k(x, x') = \exp(-\gamma||x - x'||^2)$. The different errors of the smoother g, the ode h and the trajectory $\int h(g(t))dt$ are summarized as

$$\text{Smoothing error: } \sum_{\ell=0}^{n-1} ||\mathbf{y}_\ell - g(t_\ell)||^2$$

$$\text{Gradient matching error: } \sum_{\ell=1}^{m} ||\dot{g}(t_\ell) - h(g(t_\ell))||^2$$

$$\text{Trajectory error: } \sum_{\ell=1}^{m} ||\mathbf{y}_\ell - \int_0^{t_\ell} h(g(\tau))d\tau||^2.$$

We choose the hyperparameters (γ_i, λ_i) using leave-one-out cross-validation over the dataset $\mathcal{S} = (t_i, \mathbf{y}_i)_{i=1}^n$ to minimize the empirical smoothing error. For the g, we choose the hyperparameters (γ, λ) through a grid search by minimizing the empirical trajectory error of the resulting model against the observations $\mathcal{S}$. We learn the optimal matrix C using iterative proximal descent. We solve the multiple time series model with stochastic gradient descent with batches of 10 coefficients with averaging after the first epoch, for a total of 20 epochs. For multiple time series, we set manually $\lambda = 0.1$.

6.2 ODE estimation

We apply OKODE on two dynamic models with true mechanics known, and on dataset specifically designed to represent a realistic noisy biological time series with no known underlying model.

6.2.1 Fitz-Hugh Nagumo model

The FitzHugh-Nagumo equations (Fitzhugh, 1961; Nagumo et al., 1962) approximate the spike potentials in the giant axon of squid neurons with a model

$$\dot{V} = c(V - V^3/3 + R)$$
$$\dot{R} = -1/c(V - a + bR).$$

The model describes the dependency between the voltage V across the axon membrane and the recovery variable R of the outward currents. We assign true values $(0.2, 0.2, 3)$ for the parameters (a, b, c). The FHN model is relatively simple non-linear dynamical system, however it has proven challenging to learn with numerous local optima Ramsay et al. (2007). We sample $n = 41$ observations over regularly spaced time points with added isotropic, zero-mean Gaussian noise with $\sigma^2 = 0.1$.

The Figure 6.2.1 presents the learned OKODE model with $m = 101$ parameter vectors. The four figures depict from top to bottom the (i) smoother g, (ii) the gradient $h(g(t))$ against $\dot{g}(t)$, (iii) the h against the f_true, and (iv) the estimated trajectory from h. The estimated trajectory matches the true well, but tends to underestimate the derivatives around the sharp turns of the red curve.

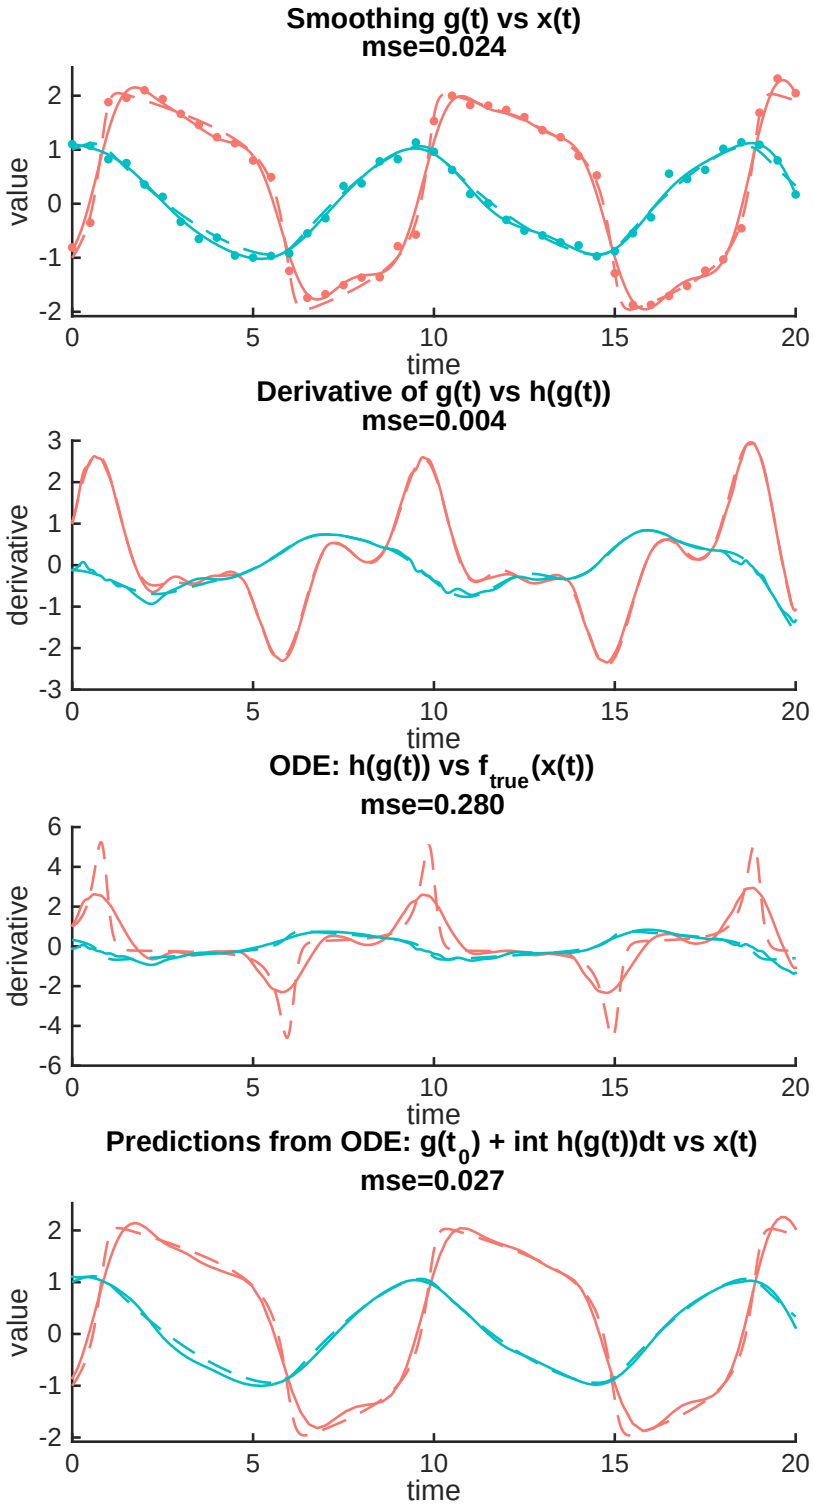

Figure 1: Learned OKODE of the FHN model.

6.2.2 Calcium model

The calcium model (Peifer and Timmer, 2007) represents the oscillations of calcium signaling in eucaryotic cells by modeling the concentrations of free calcium in cytoplasm Ca_{cyt} and in the endoplasmic reticulum Ca_{re}, and active G_α and phospholipase-C, PLP_C. The system consists of 17 parameters. We sample $n = 67$ regularly spaced observations with added zero-truncated Gaussian noise $\mathcal{N}(0, 0.1)$. We use $m = 101$ parameter vectors. 17 parameters determine the system

$$
\begin{aligned}
\dot{G_\alpha} &= k_1 + k_2 G_\alpha - k_3 P_C R_1(G_\alpha) - k_4 Ca_c R_2(G_\alpha) \\
\dot{P_C} &= k_5 G_\alpha - k_6 R_3(P_C) \\
\dot{Ca_c} &= k_7 P_C Ca_c R_4(Ca_r) + k_8 P_C + k_9 G_\alpha - k_{10} R_5(Ca_c) - k_{11} R_6(Ca_c) \\
\dot{Ca_r} &= -k_7 P_C Ca_c R_4(Ca_r) + k_1 1 R_6(Ca_c).
\end{aligned}
$$

where $R_i(x) = \frac{x}{x+Km_i}$. We sample $n = 67$ regularly spaced observations with added zero-truncated Gaussian noise $\mathcal{N}(0, 0.1)$. We use $m = 101$ parameter vectors.

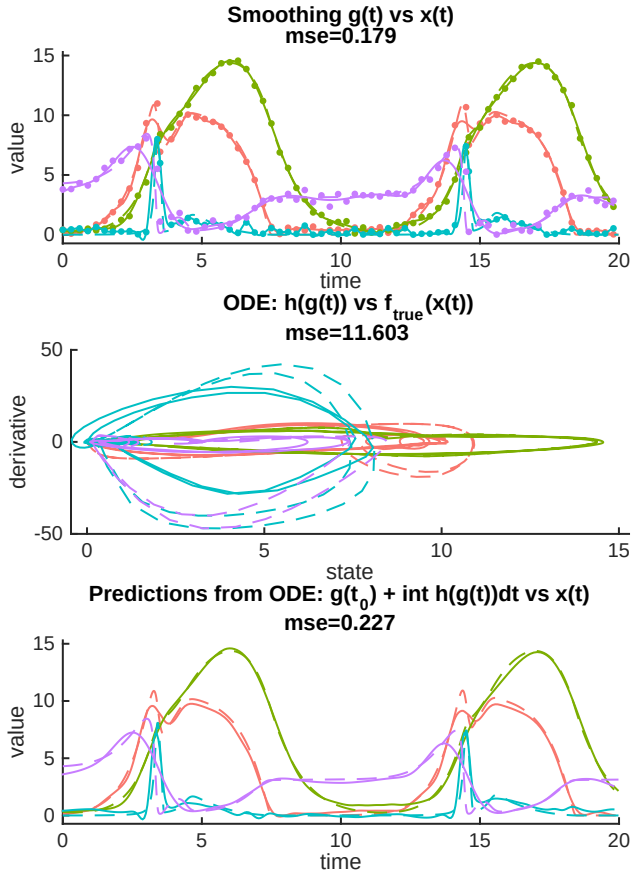

Figure 2: Learned OKODE of the calcium model.

The Figure 6.2.2 depicts the (i) smoother g, (ii) the estimated $h(x)$ and $f_{true}(x)$ over x and (iii) finally the predicted trajectory. The smoother doesn't learn the peaks of the red curve, which is reflected in the estimated trajectory.

6.2.3 DREAM3

The DREAM3 dataset consists of realistically simulated biological time-series of 10 variables with $n = 21$ noisy observations. The corresponding ODE models are deliberately unknown. We employ the DREAM3 dataset as a representative of a biological, noisy dataset with no gold standard to perform exploratory ODE modeling. We note that such data has not been applicaple to ODE analysis.

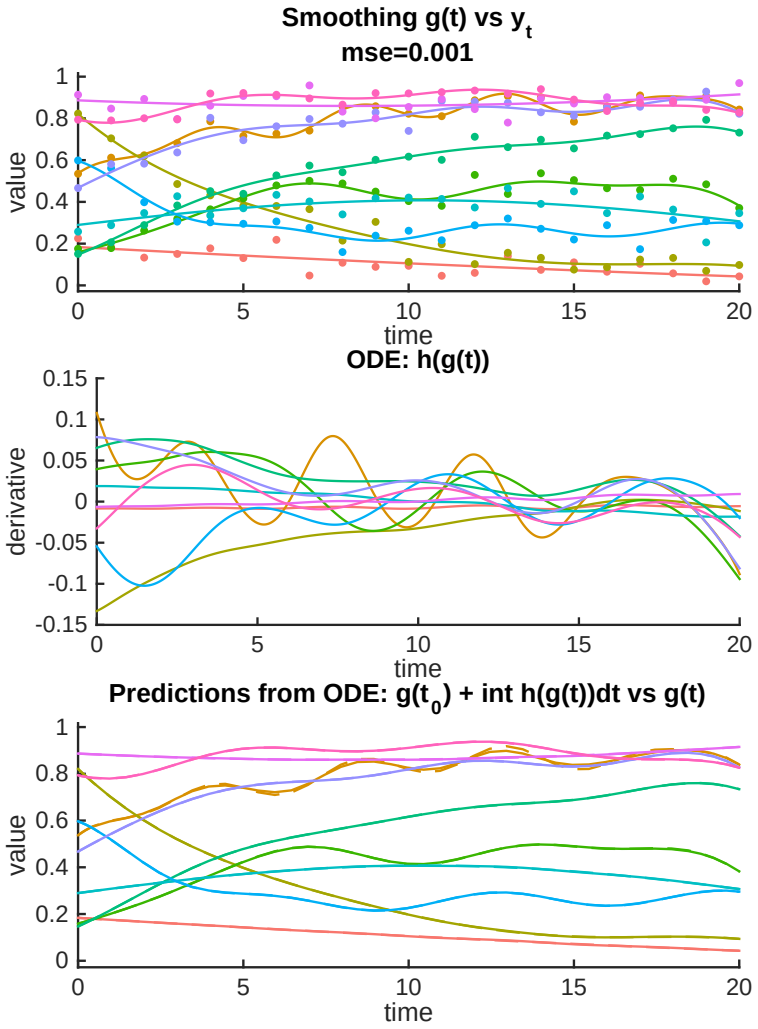

Figure 3: Learned OKODE of the DREAM3 model.

The Figure 3 shows the result of applying the proposed method to DREAM3 dataset with $m = 101$. The ODE model reconstructs the smoother g rather well, but there is considerable uncertainty in how smooth the function g should be.

6.3 Model sparsity

Figure 4 indicates true trajectory errors and levels of sparsity when learning the FHN model with different values of α and with a high number of $m = 404$ model points and $n = 41$ data points as in Section 6.2.1. Approximately half of the coefficients can be set to zero without large effect on the trajectory error. However, as an automated approach for selection of m, using sparsity is quite unrobust.

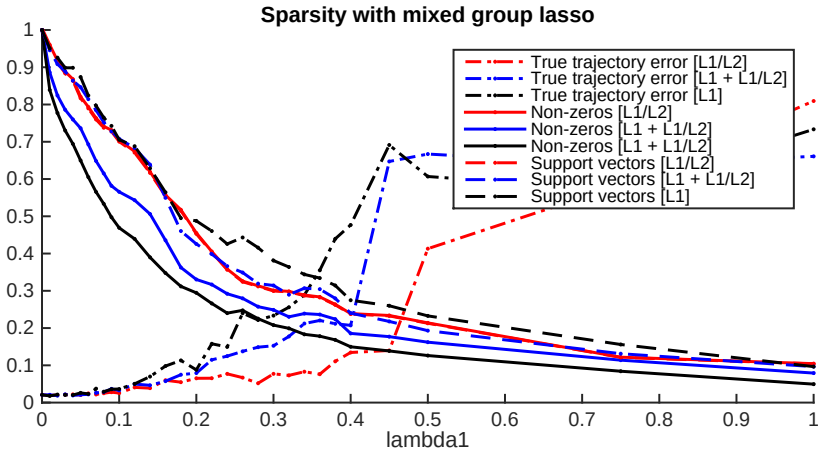

Figure 4: The tradeoff between trajectory error and sparsity under sparse norms over the FHN dataset.

6.4 Multiple initial conditions

We study learning a nonparametric model using multiple time-series on the FHN dataset. We are interested in the ability of the model to estimate accurate trajectories from arbitrary initial conditions. In Figure 5 we plot the true trajectory error of a model learned with 2 or 4 time series over the space of (V, R) values as initial values. As expected, adding more time-series improves the model's ability to generalise, and provides a more accurate ODE model.

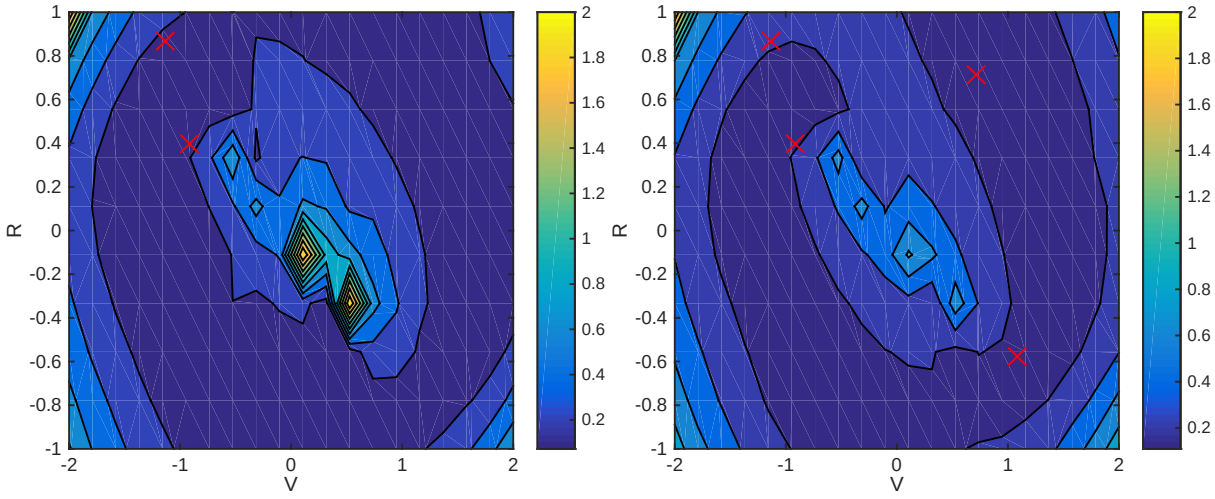

Figure 5: The generalisation error of the ODE model with a model learned from 2 or 4 time series. The red crosses indicate the initial conditions of the time-series the model was learned from.

6.5 Comparison against parametric estimation

We compare our non-parametric OKODE framework against the iterative method of Ramsay et al. (2007), and to the classic parametric parameter estimation on the FHN model with initial values $(-1, 1)$. The parametric approach has three parameters (a, b, c) to learn, and is expected to perform well as the ODE model is known. For a more realistic

comparison, we also estimate the parameters when we assume only a third order model with 14 parameters

$$\dot{V} = p_1 V + p_2 V^2 + p_3 V^3 + p_4 R + p_5 R^2 + p_6 R^3 + p_7$$
$$\dot{R} = p_8 V + p_9 V^2 + p_{10} V^3 + p_{11} R + p_{12} R^2 + p_{13} R^3 + p_{14}$$

where the true values are $p_1 = 3$, $p_3 = -1$, $p_4 = 3$, $p_8 = -1/3$, $p_{11} = p_{14} = -1/6$ and remaining values are zero. We use MATLAB's `fminsearch` and do 100 restarts from random initial values from $\mathcal{N}(0, 1)$ to estimate the best values. Table 1 highlights the true trajectory errors of the ODE. As expected, given a known ODE model, the parametric model achieves a near perfect performance. The method of Ramsay and OKODE perform well, while the parametric solver fails if the ODE model is less rigorously specified.

Table 1: Mean square errors of ODE methods on the FHN dataset.

Method	MSE
Ramsay	0.0094
Parametric, 3 free coefficients	0.0001
Parametric, 14 free coefficients	0.4925
OKODE	0.0270

7 Conclusion

We described a new framework for nonparametric ODE modeling and estimation. We showed that matrix-valued kernel-based regression were especially well appropriate to build estimates in a two-step gradient matching approach. The flexibility of penalized regression in RKHS provides a way to address realistic tasks in ODE estimation such as learning from multiple initial conditions. We show that these models can be learned as well using nonsmooth constraints with the help of proximal gradient algorithms and also discuss the relevance of sparse models. Future works concerns the study of this approach in presence of heteroscedastic noise. A Bayesian view of this approach will also be of interest (Calderhead et al. (2009); Dondelinger et al. (2013)). Finally, another important issue is the scaling up of these methods to large dynamical systems.

8 Acknowledgements

This work was supported by ANR-09-SYSC-009 and Electricité de France (Groupe Gestion Projet-Radioprotection) and Institut de Radioprotection et de Sûreté nucléaire (programme ROSIRIS).

References

M. A. Álvarez, L. Rosasco, and N. D. Lawrence. Kernels for vector-valued functions: a review. *Foundations and Trends in Machine Learning*, 4(3):195–266, 2012.

N. Aronszajn. Theory of reproducing kernels. *Transactions of the American mathematical society*, pages 337–404, 1950.

A. Beck and M. Teboulle. Gradient-based algorithms with applications to signal recovery problems. In DP Palomar and YC Eldar, editors, *Convex Optimization in Signal Processing and Communications*, pages 42–88. Cambridge press, 2010.

R. Bellman and R. Roth. The use of splines with unknown end points in the identification of systems. *J. Math. Anal. Appl.*, 34:26–33, 1971.

C. Brouard, F. d'Alché Buc, and M. Szafranski. Semi-supervised penalized output kernel regression for link prediction. In *Proc. of the 28th Int. Conf. on Machine Learning (ICML-11)*, pages 593–600, 2011.

N. J-B. Brunel. Parameter estimation of ode's via nonparametric estimators. *Electronic Journal of Statistics*, 2: 1242–1267, 2008.

B Calderhead, M Girolami, and N Lawrence. Accelerating bayesian inference over nonlinear differential equations with gaussian processes. In *NIPS*, 2009.

A. Caponnetto, C. A. Micchelli, M. Pontil, and Y. Ying. Universal multi-task kernels. *The Journal of Machine Learning Research*, 9:1615–1646, 2008.

P. L. Combettes and J.-C. Pesquet. Proximal splitting methods in signal processing. In *Fixed-point algorithms for inverse problems in science and engineering*, pages 185–212. Springer, 2011.

D. R. Cox. Tests of separate families of hypotheses. *Proc 4th Berkeley Symp Math Stat Probab*, 1:105–123, 1961.

Florence d'Alché Buc, Néhémy Lim, George Michailidis, and Yasin Senbabaoglu. Estimation of nonparametric dynamical models within Reproducing Kernel Hilbert Spaces for network inference. In *Parameter Estimation for Dynamical Systems - PEDS II, Eindhoven, June 4-6, 2012*, Eindhoven, Netherlands, June 2012. Bart Bakker, Shota Gugushvili, Chris Klaassen, Aad van der Vaart. URL `https://hal.archives-ouvertes.fr/hal-01084145`.

Francesco Dinuzzo. Learning functional dependencies with kernel methods. *Scientifica Acta*, 4(1):MS–16, 2011.

F. Dondelinger, D. Husmeier, S. Rogers, and M. Filippone. Ode parameter inference using adaptive gradient matching with gaussian processes. In *AISTATS*, volume 31 of *JMLR Proceedings*, pages 216–228. JMLR.org, 2013.

S. P. Ellner, Y. Seifu, and R. H. Smith. Fitting population dynamic models to time-series data by gradient matching. *Ecology*, 83(8):2256–2270, 2002.

T. Evgeniou and M. Pontil. Regularized multi–task learning. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 109–117. ACM, 2004.

R. Fitzhugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophys J.*, 1(6):445–66, Jul; 1961.

S. Gugushvili and C.A.J. Klaassen. Root-n-consistent parameter estimation for systems of ordinary differential equations: bypassing numerical integration via smoothing. *Bernoulli*, to appear, 2011.

M. Ha Quang and V. Bazzani, L.and Murino. A unifying framework for vector-valued manifold regularization and multi-view learning. In *Proc. of the 30th Int. Conf. on Machine Learning, ICML 2013,*, pages 100–108, 2013.

M. Hirsch, S. Smale, and Devaney. *Differential Equations, Dynamical Systems, and an Introduction to Chaos (Edition: 2)*. Elsevier Science & Technology Books, 2004.

H. Kadri, A. Rabaoui, P. Preux, E. Duflos, and A. Rakotomamonjy. Functional regularized least squares classiffication with operator-valued kernels. In *In Proc. International Conference on Machine Learning*, 2011.

N. Lim, Y. Senbabaoglu, G. Michailidis, and F. d'Alché-Buc. Okvar-boost: a novel boosting algorithm to infer nonlinear dynamics and interactions in gene regulatory networks. *Bioinformatics*, 29(11):1416–1423, 2013.

N. Lim, F. d'Alché Buc, C. Auliac, and G. Michailidis. Operator-valued Kernel-based Vector Autoregressive Models for Network Inference. hal-00872342, 2014. URL `https://hal.archives-ouvertes.fr/hal-00872342`.

C. A. Micchelli and M. Pontil. On learning vector-valued functions. *Neural Computation*, 17(1):177–204, 2005.

J. Nagumo, S Arimoto, and S Yoshizawa. An active pulse transmission line simulating nerve axon. *Proceedings of the IRE*, 50(10):2061–2070, 1962.

N. D. Pearce and M.P. Wand. Penalized splines and reproducing kernels. *The American Statistician*, 60(3), 2006.

George B Pedrick. Theory of reproducing kernels for hilbert spaces of vector valued functions. Technical report, DTIC Document, 1957.

M Peifer and J Timmer. Parameter estimation in ordinary differential equations for biochemical processes using the method of multiple shooting. *Systems Biology, IET*, 1(2):78–88, 2007.

J.O. Ramsay, G. Hooker, J. Cao, and D. Campbell. Parameter estimation for differential equations: A generalized smoothing approach. *Journal of the Royal Statistical Society (B)*, 69:741–796, 2007.

E. Richard, P.-A. Savalle, and N. Vayatis. Estimation of simultaneously sparse and low rank matrices. In John Langford and Joelle Pineau, editors, *ICML-2012*, pages 1351–1358, New York, NY, USA, July 2012. Omnipress. ISBN 978-1-4503-1285-1.

E. Senkene and A. Tempel'man. Hilbert spaces of operator-valued functions. *Lithuanian Mathematical Journal*, 13 (4):665–670, 1973.

J. M. Varah. A spline least squares method for numerical parameter estimation in differential equations. *SIAM J.sci. Stat. Comput.*, 3(1):28–46, 1982.

QH Vuong. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, 57(2):307–333, 1989. doi: 10.2307/1912557.

G. Wahba. *Spline model for observational data*. Philadelphia, Society for Industrial and Applied Mathematics, 1990.

M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *J. of the Royal Statistical Society: Series B*, 68(1):49–67, 2006.